

Convergence Rates for Mixture-of-Experts

Eduardo F. Mendes

Wenxin Jiang

Department of Statistics, Northwestern University

November 2, 2011

Abstract

In mixtures-of-experts (ME) model, where a number of submodels (experts) are combined, there have been two longstanding problems: (i) how many experts should be chosen, given the size of the training data? (ii) given the total number of parameters, is it better to use a few very complex experts, or is it better to combine many simple experts? In this paper, we try to provide some insights to these problems through a theoretic study on a ME structure where m experts are mixed, with each expert being related to a polynomial regression model of order k . We study the convergence rate of the maximum likelihood estimator (MLE), in terms of how fast the Kullback-Leibler divergence of the estimated density converges to the true density, when the sample size n increases. The convergence rate is found to be dependent on both m and k , and certain choices of m and k are found to produce optimal convergence rates. Therefore, these results shed light on the two aforementioned important problems: on how to choose m , and on how m and k should be compromised, for achieving good convergence rates.

Keywords: Convergence Rate, Approximation Rate, Nonparametric Regression, Exponential Family, Hierarchical Mixture-of-Experts, Mixture-of-Experts, Maximum Likelihood estimation

1 Introduction

Mixture-of-experts models (ME) [Jacobs et al., 1991] and hierarchical mixture-of-experts models (HME) [Jordan and Jacobs, 1994] are powerful tools for estimating the den-

sity of a random variable Y conditional on a known set of covariates X . The idea is to “divide-and-conquer”. We first divide the covariate space into *subspaces*, then approximate each *subspace* by an adequate model and, finally, weigh by the probability that X falls in each *subspace*. Additionally, it can be seen as a generalization of the classical mixture-of-models, whose weights are constant across the covariate space. Mixture-of-experts have been widely used on a variety of fields including image recognition and classification, medicine, audio classification and finance. Such flexibility have also inspired a series of distinct models including [Wood et al. \[2002\]](#), [Carvalho and Tanner \[2005a\]](#), [Geweke and Keane \[2007\]](#), [Wood et al. \[2008\]](#), [Villani et al. \[2009\]](#), [Young and Hunter \[2010\]](#) and [Wood et al. \[2011\]](#), among many others.

We consider a framework similar to [Jiang and Tanner \[1999a\]](#) among others. Assume each expert is in a one-parameter exponential family with mean $\varphi(h_k)$, where h_k is a k^{th} -degree polynomial on the conditioning variables X (hence a linear function of the parameters) and $\varphi(\cdot)$ is the inverse link function. In other words, each expert is a Generalized Linear Model on an one-dimensional exponential family (GLM1). We allow the target density to be in the same family of distributions, but with conditional mean $\varphi(h)$ with $h \in \mathcal{W}_{\alpha, K_0}^\infty$, a Sobolev class with α derivatives. Some examples of target densities include the Poisson, binomial, Bernoulli and exponential distributions with unknown mean. Normal, gamma and beta distributions also fall in this class if the dispersion parameter is known.

One might be reluctant to use (H)ME models with polynomial experts since it leads to more and more complex models as the degree k of the polynomials increases. The discussion whether is better to mixture many simple models or fewer more complex models is not new in the literature of mixture-of-experts. Earlier in the literature, [Jacobs et al. \[1991\]](#) and [Peng et al. \[1996\]](#) proposed mixtures of many simple models; more recently, [Wood et al. \[2002\]](#) and [Villani et al. \[2009\]](#) considered using only a few complex models. [Celeux et al. \[2000\]](#) and [Geweke \[2007\]](#) advocate for mixing fewer complex models, claiming that mixture models can be very difficult to estimate and interpret. We justify the use of such models through the approximation and estimation errors. We illustrate that might be a gain in a small increase of k compared to the linear model $k = 1$ but the number of parameters increases exponentially as k increases.

Therefore, a balance between the complexity of the model and the number of experts is required for achieving better error bounds.

This work extends [Jiang and Tanner \[1999a\]](#) in few directions. We show that, by including polynomial terms, one is able to improve the approximation rate on sufficiently smooth classes. This rate is sharp for the piecewise polynomial approximation as shown in [Windlund \[1977\]](#). Moreover, we contribute to the literature by providing rates of convergence of the maximum likelihood estimator to the true density. We emphasize that such rates have never been developed for this class of models and the method used can be straightforwardly generalized to more general classes of mixture of experts. Convergence of the estimated density function to the true density and parametric convergence of the quasi-maximum likelihood estimator to the pseudo-true parameter vector are also obtained.

We found that, under slightly weaker conditions than [Jiang and Tanner \[1999a\]](#), the approximation rate in Kullback-Leibler divergence is uniformly bounded by $c \times m^{-2[\alpha \wedge (k+1)]/s}$, where c is some constant not depending on k or m , and s the number of independent variables. This is a generalization of the rate found in [Jiang and Tanner \[1999a\]](#) who assume $\alpha = 2$ and $k = 1$. The convergence rate of the maximum likelihood estimator to the true density is $O_p(m^{-2[\alpha \wedge (k+1)]/s} + (mJ_k + v_m)n^{-1} \log n)$, where J_k is the total number of parameters in each polynomial (typically $k + s$ choose k), and v_m is the number of parameters in the weight functions. To show the previous results we do not assume identifiability of the model as it is natural for mixture-of-experts to be unidentifiable under permutation of the experts. If we further assume identifiability [[Jiang and Tanner, 1999a](#), [Mendes et al., 2006](#)], and that the likelihood function has a unique maximizer, we are able to remove the “ $\log n$ ” term in the convergence rate. Optimal nonparametric rates of convergence can be attained if $k = \alpha - 1$ and $m = O(n^{s/(2\alpha+s)})$ [[Chen, 2006](#), [Stone, 1980, 1985](#)].

[Zeevi et al. \[1998\]](#) show approximation in the L^p norm and estimation error for the conditional expectation of the ME with generalized linear experts. [Jiang and Tanner \[1999a\]](#) show consistency and approximation rates for the HME with generalized linear model as experts and a general specification for the gating functions. They consider the target density to belong to the exponential family with one parameter. Their approximation rate of the Kullback-Leibler divergence between the target density and the model

is $O(1/m^{4/s})$, where m the number of experts and s the number of covariates. [Norets \[2010\]](#) show the approximation rate for the mixture of Gaussian experts where both the variance and the mean can be nonlinear and the weights are given by multinomial logistic functions. He considers the target density to be a smooth continuous function and the dependent variable Y to be continuous and satisfy some moment conditions. His approximation rate is $O(1/m^{s+2+1/(q-2)+\varepsilon})$, where Y is assumed to have at least q moments and ε is a small number. Despite these findings, there are no convergence rates yet for the maximum likelihood estimator of mixture-of-experts type of models in the literature.

By studying the convergence rates in this paper, we will be able to shed light on two long-standing problems in ME: *(i) How to choose the number of experts m for a given sample size n ?* *(ii) Is it better to mix many simple experts or to mix a few complex experts?* None of the works discussed above directly address these questions. Our study of a ME structure mixing m of the k th order polynomial submodels is particularly useful in studying problem (i), which cannot be studied in the framework of [\[Jiang and Tanner, 1999a\]](#), for example, who have restricted to the special case $k = 1$.

Throughout the paper we use the following notation. Let $x = (x_1, \dots, x_s) \in S$ and $h(x) : S \rightarrow \mathbb{R}$ an λ denote some measure. For any finite vector x we use $|x| = \sum_{j=1}^s |x_j|$ and $|x|_p = \left(\sum_{j=1}^p |x_j|^p\right)^{1/p}$, for $p \in [1, \infty)$, if $p = \infty$ we take $|x|_\infty = \sup_{j=1, \dots, s} |x_j|$. For some function $h(x)$ and measure λ we denote $\|h\|_{p,S} = (\int_S |h|^p d\lambda)^{1/p}$, for $p \in [1, \infty)$, and for $p = \infty$ we have $\|h\|_{\infty,S} = \text{ess sup}_{x \in S} |h(x)|$.

The remainder of the paper is organized as follows. In the next section we introduce the target density and mixture of experts models. We also demonstrate that the quasi-maximum likelihood estimator converges to the pseudo-true parameter vector. Section 3 establishes the main results of the paper: approximation rate, convergence rate and non-parametric consistency. Section 4 discusses model specification and the tradeoff that we unveil between the number of experts and the degree of the polynomials. In the concluding remarks we compare our results with [Jiang and Tanner \[1999a\]](#) and provide direction for future research. The appendix collects technical details of the paper and a deeper treatment on how to bound the estimation error.

2 Preliminaries

In this section we introduce the target class of density, mixture-of-experts model with GLM1 experts and the estimation algorithm.

2.1 Target density

Consider a sequence of random vectors $\{(X'_i, Y_i)'\}_{i=1}^n$ defined on $((\Omega \times A)^n, \mathcal{B}_{(\Omega \times A)^n}, P_{xy}^n)$ where $X \in \Omega \subset \mathbb{R}^s$, $Y \in A \subseteq \mathbb{R}$ and \mathcal{B}_S is the Borel σ -algebra generated by the set S . We assume that P_{xy} has a density $p_{xy} = p_{y|x}p_x$ with respect to some measure λ . More precisely, we assume that p_x is known and $p_{y|x}$ is member of an one-dimensional exponential family, i.e.

$$p_{y|x} = \exp \{ya(h(x)) + b(h(x)) + c(y)\}, \quad (1)$$

where $a(\cdot)$ and $b(\cdot)$ are known three times continuously differentiable functions, with first derivative bounded away from zero and $a(\cdot)$ has a non-negative second derivative; $c(\cdot)$ is a known measurable function of Y . The function $h(\cdot)$ is a member of $\mathcal{W}_{\alpha, K_0}^\infty(\Omega)$, a Sobolev class of order α ¹. Throughout the paper denote by $\Pi(\mathcal{W}_{\alpha, K_0}^\infty)$ the class of density functions $p_{xy} = p_{y|x}p_x$.

The one-parameter exponential family of distributions includes the Bernoulli, exponential, Poisson and binomial distributions, it also includes the Gaussian, gamma and Weibull distributions if the dispersion parameter is known. It is possible to extend the results to the case where the dispersion parameter is unknown, but defined in a compact subset bounded away from zero. In this work we focus only in the one-parameter case.

Some properties of the one parameter exponential family are : (i) conditional on $X = x$, the moment generating function of y exists in a neighborhood of the origin implying that moments of all orders exist; (ii) for each positive integer j , $\mu_{(j)}(h) = \int_A y^j \exp[a(h)y + b(h) + c(y)]d\lambda$ is a differentiable function of h ; and (iii) the first conditional moment $\mu_{(1)}(h) = -\dot{b}(h)/\dot{a}(h) = \varphi(h)$, where $\dot{a}(h)$ and $\dot{b}(h)$ are the first derivatives of $a(h)$ and $b(h)$ respectively, and $\varphi(\cdot)$ is called the inverse link function.

¹ Suppose $1 \leq p \leq \infty$ and $\alpha > 0$ is an integer. We define $\mathcal{W}_{p, K_0}^\alpha(\Omega)$ as the collection of measurable functions h with all distributional derivatives $D^r f$, $|r| \leq \alpha$, on $L^p(\Omega)$, i.e. $\|D^r h\|_{p, \Omega} \leq K_0$. Here $D^r = \partial^{|r|}/(\partial^{r_1} x_1 \dots \partial^{r_s} x_s)$ and $|r| = r_1 + \dots + r_s$ for $r = (r_1, \dots, r_s)$.

See [Lehmann \[1991\]](#) and [McCullagh and Nelder \[1989\]](#) for more results about the exponential family of distributions.

2.2 Mixture-of-experts model

The mixture-of-experts model with GLM1 experts is defined as:

$$\begin{aligned} f_{m,k}(x, y; \zeta) &= \sum_{j=1}^m g_j(x; \nu) \pi(h_k(x; \theta_j), y) \cdot p_x \\ &= \sum_{j=1}^m g_j(x; \nu) \exp\{ya(h_k(x; \theta_j)) + b(h_k(x; \theta_j)) + c(y)\} \cdot p_x, \quad (2) \end{aligned}$$

where the functions $g_j \geq 0$ and $\sum_{j=1}^m g_j = 1$ with parameters $\nu \in V_m \subset \ell^2(\mathbb{R}^{v_m})^2$, v_m denoting the dimension of ν . The functions $h_k(x; \theta_j)$ are k^{th} -degree polynomials on Ω with parameter vector $\theta_j \in \Theta_k \subset \ell^2(\mathbb{R}^{J_k})$, J_k denoting the dimension of θ_j ; write the vector of parameters of all experts as $\theta = (\theta'_1, \dots, \theta'_m)'$ defined on $\Theta_{mk} \equiv \Theta_k^m$. The parameter vector of the model is $\zeta = (\nu', \theta')'$ and is defined on $V_m \times \Theta_{mk}$, a subset of $\mathbb{R}^{v_m \times m J_k}$. Throughout the paper we denote by $\mathcal{F}_{m,k}$ the class of (approximant) densities $f_{m,k}$.

To derive consistency and convergence rates, one need to impose some restrictions on the functions π and g_j to avoid abnormal cases. This condition is not restrictive and is satisfied by the multinomial logistic weight functions (g 's) and the Bernoulli, binomial, Poisson and exponential experts, among many other classes of distributions and weight functions.

Assumption 1. *There exist functions $c_g(x) = (c_g^{(1)}(x), \dots, c_g^{(v_m)}(x))'$ and $F(x, y) = (F^{(1)}(x, y), \dots, F^{(J_k)}(x, y))'$ with $\mathbb{E}[c_g(X)'c_g(X)] < \infty$ and $\mathbb{E}[F(X, Y)'F(X, Y)] < \infty$, such that the vector-function $g(x; \nu) = (g_1(x; \nu), \dots, g_m(x; \nu))'$ satisfy*

$$\sup_{\nu \in V_m} \frac{\partial \log g(x; \nu)}{\partial \nu_i} \leq c_g^{(i)}(x);$$

and each expert $\pi(h_k(x; \theta_j), y)$ satisfy

$$\sup_{\theta \in \Theta_k} \frac{\partial \pi(h_k(x; \theta_j), y)}{\partial \theta_{ji}} \leq F^{(i)}(x, y), \quad \text{for each } 1 \leq j \leq m.$$

²We denote $\ell^2(\mathbb{R}^k) \equiv \{x \in \mathbb{R}^k : \sum_{j=1}^k x_j^2 < \infty\}$.

2.3 Maximum likelihood estimation and the EM algorithm

2.1 Maximum likelihood estimation

We consider the maximum likelihood method of estimation. We want to find the parameter vector $\hat{\zeta}_n = (\hat{\nu}'_n, \hat{\theta}'_n)'$ that maximizes

$$L_n(\zeta) = n^{-1} \sum_{i=1}^n \log \{f_{m,k}(X_i, Y_i; \zeta)/\varphi_0(X_i, Y_i)\}, \quad (3)$$

where $\varphi_0(X, Y) = \exp(c(Y))p_x(X)$. That is,

$$\hat{\zeta}_n = \arg \max_{\zeta \in V_m \times \Theta_{mk}} L_n(\zeta). \quad (4)$$

The maximum likelihood estimator is not necessarily unique. In general, mixture-of-experts models are not identifiable under permutation of the experts. To circumvent this issue one must impose restrictions on the experts and the weighting (or the parameter vector of the model), as shown in [Jiang and Tanner \[1999b\]](#).

Define the Kullback-Leibler (KL) divergence between p_{xy} and $f_{m,k}$ as

$$KL(p_{xy}, f_{m,k}) = \int_{\Omega} \int_A \log \frac{p_{xy}}{f_{m,k}} dP_{y|x} dP_x. \quad (5)$$

The log-likelihood function in (3) converges to its expectation with probability one as the number of observations increases. Therefore, in the limit, the minimizer $\hat{f}_{m,k}$ of (3) (indexed by $\hat{\zeta}_n$) also minimizes the Kullback-Leibler divergence between the true density and the estimated density.

In this work only consider i.i.d. observations but is straightforward to extend the results to more general data generating processes. Next assumption formalizes it.

Assumption 2 (Data Generating Process). *The sequence $(X_i, Y_i)_{i=1}^n$, $n = 1, 2, \dots$ is an independent and identically distributed sequence of random vectors with common distribution P_{xy} .*

Next results ensures the existence of such estimator.

Theorem 2.1 (Existence). *For a sequence $\{(V_m \times \Theta_{mk})_n\}$ of compact subsets of $V_m \times \Theta_{mk}$, $n = 1, 2, \dots$, there exists a $\mathcal{B}(\Omega \times A)$ -measurable function $\hat{\zeta}_n : \Omega \times A \rightarrow (V_m \times \Theta_{mk})_n$, satisfying equation (4) P_{xy} -almost surely.*

We demonstrate that under the classical assumptions, such as identifiability and unique maximizer, the maximum likelihood estimator $\hat{\zeta}$ consistently estimate the best model in the class \mathcal{F}_{mk} indexed by ζ^* , i.e. the maximum likelihood estimator $\hat{\zeta}$ converges almost surely to ζ^* . It can be shown that the convergence results also hold for the ergodic case if we assume that $(\log f_{m,k}(X_i, Y_i; \zeta))_{i=1}^n$ is ergodic. However, simpler conditions to ensure ergodicity of the likelihood function are not trivial and hence out of the scope of this paper.

Assumption 3 (Identifiability). *For any distinct ζ_1 and ζ_2 in $V_m \times \Theta_{mk}$, for almost every $(x, y) \in \Omega \times A$,*

$$f_{m,k}(x, y; \zeta_1) \neq f_{m,k}(x, y; \zeta_2)$$

[Jiang and Tanner \[1999b\]](#) find sufficient conditions for identifiability of the parameter vector for the HME with one layer, while [Mendes et al. \[2006\]](#) for a binary tree structure. Both cases can be adapted to more general specifications. Although one can show consistency to a set, we adopt a more traditional approach requiring identifiability of the parameter vector.

Assumption 4 (Unique Maximizer). *Let $\zeta = (\nu', \theta')'$ and ζ^* the argument that maximizes $\mathbb{E} \log f_{m,k}$ over $V_m \times \Theta_{mk}$. Then*

$$\det \left(\mathbb{E} \frac{\partial^2}{\partial \zeta \partial \zeta'} \log f_{m,k} |_{\zeta=\zeta^*} \right) \neq 0 \quad (6)$$

This assumption follows from a second order Taylor expansion of the expected likelihood around the parameter vector that maximizes (5), denoted ζ^* . We require the Hessian to be invertible at ζ^* . The requirement for an identifiable unique maximizer is only technical in a sense that the objective function is not allowed to become too flat around the maximum (For more discussion on this topic see [Bates and White \[1985\]](#), pg 156, and [White \[1996\]](#) chapter 3). A similar assumption was made in the series of papers from [Carvalho and Tanner \[2005a,b, 2006, 2007\]](#) and [Zeevi et al. \[1998\]](#) and is an usual assumption in the estimation of misspecified models.

Theorem 2.2 (Parametric consistency of misspecified models). *Under Assumptions 1, 2, 3, and 4, the maximum likelihood estimate $\hat{\zeta} \rightarrow \zeta^*$ as $n \rightarrow \infty$ P_{xy} -a.s.*

Huerta et al. [2003] and the series of papers by Carvalho and Tanner [2005a,b, 2006, 2007] derive similar results for time series processes.

2.2 The EM algorithm

It is often easier to maximize the *complete likelihood function* of a (H)ME instead of (3) (see Jordan and Jacobs [1994], Xu and Jordan [1996] and Yang and Ma [2011]). Let $z'_i = (z_{i1}, \dots, z_{im})$ denote a binary vector with $z_{ij} = 1$ if the observation (x_i, y_i) is generated by the expert j (i.e. $\pi(h_k(\cdot, \theta_j), \cdot)$). We assume z_i has a multinomial distribution with parameters $\tau'_i = (\tau_{i1}, \dots, \tau_{im})$. The complete log-likelihood function is given by

$$l_n^c(\kappa) = \sum_{i=1}^n \sum_{j=1}^m z_{ij} (\log g_i(x_i, \nu) + \log \pi(h_k(x_i; \theta_j), y_i) - \log \varphi_0(x_i, y_i)), \quad (7)$$

where $\kappa = (\nu', \theta', \tau')'$.

We can estimate this model using the expectation-maximization (EM) algorithm put forward by Dempster et al. [1977]. Let $\kappa^{(l)} = (\nu^{(l)}, \theta^{(l)}, \tau^{(l)})$ denote the parameter estimates at the l th iteration and define $q(\kappa; \kappa^{(l)}) = \mathbb{E}(l_n^c | x, y; \kappa^{(l)})$. In the E-step, we obtain $q(\kappa, \kappa^{(l)})$ by replacing z_{ij} with its expectation

$$\tau_{ij}^{(l)} = \frac{g_j(x_i, \nu^{(l)}) \pi(h_k(x_i; \theta_j^{(l)}), y_i)}{\sum_{j=1}^m g_j(x_i; \nu^{(l)}) \pi(h_k(x_i; \theta_j^{(l)}), y_i)}.$$

In the M-step we maximize $q(\kappa; \kappa^{(l)})$ with respect to ν and θ . The problem simplifies to find the parameters $\nu^{(l+1)}$ that maximize

$$q(\nu; \kappa^{(l)}) = \sum_{i=1}^n \tau_{ij}^{(l)} \log g_j(x_i; \nu), \quad (8)$$

and to find the parameters $\theta^{(l+1)}$ we have to maximize

$$q(\theta; \kappa^{(l)}) = \sum_{i=1}^n \sum_{j=1}^k \tau_{ij}^{(l)} [y_i a(h_k(x_i; \theta_j)) + b(h_k(x_i; \theta_j))]. \quad (9)$$

3 Main results

In this section we present the main results of the paper. Write the KL-divergence as follows:

$$KL(p_{xy}, f_{m,k}) = KL(p_{xy}, f_{m,k}^*) + \mathbb{E} \left[\log \frac{f_{m,k}^*}{f_{m,k}} \right], \quad (10)$$

where $f_{m,k}^*$ is the minimizer of the minimizer of $KL(p_{xy}, f_{m,k})$ on $\mathcal{F}_{m,k}$. The first term in the right-hand side is the **approximation error** and the second term is the **estimation error**. The approximation error measures “how well” an element of $\mathcal{F}_{m,k}$ approximates p_{xy} , and approaches zero as m increases. The estimation error measures “how far” is the estimated model from the best approximant in the class. Our goal is to find bounds for both approximation and estimation errors and combine these results to find the convergence rate of the maximum likelihood estimator.

3.1 Approximation rate

We follow [Jiang and Tanner \[1999a\]](#) to bound the approximation error. Define the *upper divergence* between $p \in \Pi(\mathcal{W}_{\alpha,K_0}^\infty)$ and $f_{m,k} \in \mathcal{F}_{m,k}$ as

$$\mathcal{D}(p, f_{m,k}) = \int_{\Omega} \sum_{j=1}^m g_j(x, \nu) (h_k(x; \theta_j) - h(x))^2 dP_x. \quad (11)$$

We can use the upper divergence to bound the KL-divergence.

Lemma 3.1. *Let $p \in \Pi(\mathcal{W}_{\alpha,K_0}^\infty)$ and $f_{m,k} \in \mathcal{F}_{m,k}$. If $\text{ess sup}_{x \in \Omega} |h(x)| < \infty$,*

$$KL(p, f_{m,k}) \leq M_\infty \mathcal{D}(p, f_{m,k})$$

where $M_\infty \geq (1/2) \text{ess sup}_{x \in \Omega} [|\varphi(h(x))| \cdot |\ddot{a}(h(x))| + |\ddot{b}(h(x))|]$.

This lemma will be used to bound uniformly the approximation rate of the family of functions $\mathcal{F}_{m,k}$.

Before presenting the main conditions, we shall introduce some key concepts.

Definition 3.1 (Fine partition). *For $m = 1, 2, \dots$, let $Q^m = \{Q_j^m\}_{j=1}^{r_m}$ be a partition of Ω . If $m \rightarrow \infty$ and if for all $x_1, x_2 \in Q_j^m$, $\max_{1 \leq i \leq s} |(x_1 - x_2)_i| \leq c_0/r_m^{1/s}$, for some constant c_0 independent of x_1, x_2, m or j . Then $\{Q^m, m = 1, 2, \dots\}$ is called a sequence of fine partitions with cardinality r_m and bounding constant c_0 .*

Here we use some abuse of notation by using m as an index of the collection of partitions of Ω . However, this abuse of notation is justified because m is an increasing sequence and the collection of partitions depends on an increasing function of m . The next definition will be useful later to bound the “growth rate” of the model and is useful to deal with hierarchical mixture of experts (see [Jiang and Tanner \[1999a\]](#)).

Definition 3.2 (Subgeometric). *A sequence of numbers a_j is called sub-geometric with rate bounded by M_1 if $a_j \in \mathbb{N}$, $a_j \rightarrow \infty$ as $j \rightarrow \infty$, and $1 < |a_{j+1}/a_j| < M_1$ for all $j = 1, 2, \dots$ and for some finite constant M_1 .*

The key idea behind find the approximation rates, is to control the approximation rate inside each fine partition of the space. More precisely, bound the approximation inside the “worst” (more difficult to approximate) partition. We need the following assumption.

Assumption 5. *There exists a fine partition Q^m of Ω , with bounding constant c_0 and cardinality sequence r_m , $m = 1, 2, \dots$, such that $\{r_m\}$ is sub-geometric with rate bounded by some constant M_1 , and there exists a constant $c_1 > 0$, and a parameter vector $\nu_{c_1} \in V_m$ such that*

$$\max_{1 \leq j \leq r_m} \|g_j(\cdot; \nu_{c_1}) - I_{Q_j^m}(\cdot)\|_{1,\lambda} \leq \frac{c_1}{r_m}. \quad (12)$$

This assumption is similar, but weaker than, the one employed in [Jiang and Tanner \[1999a\]](#) and requires that the vector $g = (g_1, \dots, g_{r_m})$ approximates the vector of characteristic functions $(I_{Q_1^m}, \dots, I_{Q_{r_m}^m})$ at a rate not slower then $O(r_m)$.

The notation r_m is introduced to deal with the hierarchical mixture of experts structure. To allow more flexibility define r_m as the maximum number of experts the structure can hold, e.g. a binary tree with $l = 1, 2, \dots$ layers has at most 2^l experts, and if we increase the number of layers by one, the actual number of experts is somewhere between 2^l and $2^{l+1} - 1$ (here we are assuming the tree is balanced without loss of generality). If we denote this class of models by $\mathcal{F}_{r_m,k}^*$, then $\mathcal{F}_{r_m,k}^* \subseteq \mathcal{F}_{m,k}^* \subset \mathcal{F}_{r_{m+1},k}^*$. The sub-geometric assumption ensures that $r_m \leq m < r_{m+1}$, where m is the actual number of experts in the model.

Theorem 3.1 (Approximation rate). *Let $p \in \Pi(\mathcal{W}_{\alpha, K_0}^\infty)$ and $f_{m,k} \in \mathcal{F}_{m,k}^*$. If assumption 5 holds, then*

$$\sup_p \inf_{f_{m,k}} KL(p, f_{m,k}) \leq \frac{c}{m^{2(\alpha \wedge (k+1))/s}} \quad (13)$$

for some constant c not depending on m or k .

This result is a generalization of [Jiang and Tanner \[1999a\]](#) in two directions. First we allow the target function to be in a Sobolev class with α derivatives; second, we consider a polynomial approximation to the target function in each experts (in fact, their result is a special case when $\alpha = 2$ and $k = 1$). This generalization enables us to address the important problem: whether it is better to mix many simple experts or to mix a few complex experts. The result also holds under more general specifications of densities/experts. In the case we also have a dispersion parameter to estimate, we just have to modify the lemma 3.1 accordingly and the same result holds.

This rate also agrees with the optimal approximation rate of functions on $\mathcal{W}_{\alpha, K_0}^\infty$ by piecewise polynomials [[Windlund, 1977](#)]. One can see that, under assumption 5, it is exactly what we are doing. Therefore this approximation rate is sharp.

3.2 Convergence rate

In this section we deduce the convergence rate for the mixture-of-experts model. Equation (10) gives us an expansion of the KL divergence in terms of the approximation and estimation errors. In the previous section we found a bound for the approximation error, in this section we will find the estimation error and combine with the approximation error to find the rate of convergence.

The estimation error is the “how far” is the estimated function from the best approximant in the class. We will demonstrate that the estimation error in (10) is $O_p((mJ_k + v_m)(\log n/n))$. We also show that by combining this result with the approximation rate it is possible to achieve a convergence rate of $O_p((\log n/n)^{2\tau/(2\tau+s)})$, with $\tau = \alpha \wedge (k+1)$, which is close to the optimal nonparametric rate if $\tau = \alpha$. Moreover, if there is an unique identifiable maximizer to the likelihood problem (assumptions 3 and 4), we are able to remove the “ $\log n$ ” term and achieve a better convergence rate, possibly optimal if $\tau = \alpha$.

The next theorem summarizes the convergence rate of the maximum likelihood estimator $\hat{f}_{m,k}$ with respect the KL divergence between the true density p_{xy} and the estimated density.

Theorem 3.2 (Convergence Rate). *Let $p_{xy} \in \Pi(\mathcal{W}_{\alpha,K_0}^\infty)$ and $\hat{f}_{m,k}$ denote its maximum likelihood estimator on $\mathcal{F}_{m,k}$. Let m be allowed to increase such that $m \rightarrow \infty$ and $m(\log n/n) \rightarrow 0$ as n and m increase. Under Assumptions 1, 2 and 5,*

$$KL(p_{xy}, \hat{f}_{m,k}) = O_p \left(\frac{1}{m^{2\tau/s}} + (mJ_k + v_m) \frac{\log n}{n} \right), \quad (14)$$

where $\tau = \alpha \wedge (k+1)$. In particular, if we assume $v_m = O(m)$, and let m be proportional to $(n/\log n)^{s/(2\tau+s)}$ then

$$KL(p_{xy}, \hat{f}_{m,k}) = O_p \left(\left(\frac{\log n}{n} \right)^{\frac{2\tau}{2\tau+s}} \right). \quad (15)$$

Although the previous result is derived for the i.i.d. case, the result also holds for more general data generating process. In this result we use (through van der Geer [2000]), an uniform probability inequality for i.i.d. processes to derive theorem A.1, but the same result can be obtained by using uniform inequalities for more general processes. This convergence rate is close to the optimal rate found in the sieves literature if $\tau = \alpha$, see for instance Stone [1980] and Barron and Sheu [1991].

To derive this rate we do not assume that there is an unique identifiable maximizer $f_{m,k}^*$; in fact, we assume $f_{m,k}^*$ is any of such maximizers. The price to pay for such generality is the inclusion of the “ $\log n$ ” term in the convergence rates. If we assume $f_{m,k}^*$ is unique and uniquely identified by a parameter vector ζ^* , we can explore the localization property of theorem A.1. More precisely, we can explore the fact that we are only interested in the behavior of the empirical process around a neighborhood of $f_{m,k}^*$. Under such conditions and assuming $\tau = \alpha$, we are able to achieve the optimal convergence rate in the sieves literature [Barron and Sheu, 1991, Stone, 1980].

Theorem 3.3 (Optimal Convergence Rate). *Let $p_{xy} \in \Pi(\mathcal{W}_{\alpha,K_0}^\infty)$ and $\hat{f}_{m,k}$ denote its maximum likelihood estimator on $\mathcal{F}_{m,k}$. Let m be allowed to increase such that $m \rightarrow \infty$ and $m/n \rightarrow 0$ as n and m increase. Under Assumptions 1–5,*

$$KL(p_{xy}, \hat{f}_{m,k}) = O_p \left(\frac{1}{m^{2\tau/s}} + \frac{(mJ_k + v_m)}{n} \right), \quad (16)$$

where $\tau = \alpha \wedge (k + 1)$. In particular, if we assume $v_m = O(m)$, and let m be proportional to $n^{s/(2\tau+s)}$ then

$$KL(p_{xy}, \hat{f}_{m,k}) = O_p \left(n^{-\frac{2\tau}{2\tau+s}} \right). \quad (17)$$

The same result follows for more general data generating processes and the same considerations after theorem 3.2 hold.

By imposing there exist an unique maximum we are able to remove the $\log n$ term and recover the optimal convergence rate for sieves estimates found in the literature.

3.3 Consistency

Now we apply the previous results to show the maximum likelihood estimator is consistent, i.e. the KL divergence between the true density and the estimated model approaches zero as the sample size n , and the index of the approximation class m goes to infinity. Here we show consistency essentially by using the previous results.

Corollary 3.1 (Consistency). *Let $p_{xy} \in \Pi(\mathcal{W}_{\alpha,K_0}^\infty)$ and $\hat{f}_{m,k}$ denote its maximum likelihood estimator on $\mathcal{F}_{m,k}$. Allow $m \rightarrow \infty$ and $m(\log n/n) \rightarrow 0$ as n and m increase. Under Assumptions 1, 2 and 5, $KL(p_{xy}, \hat{f}_{m,k}) \rightarrow 0$ as n and m increase.*

4 Effects of m and k

We consider a framework similar to [Jiang and Tanner \[1999a\]](#), but one is allowed to mix m GLM1 experts whose terms are polynomials on the variables, as opposed to $k = 1$. We also assume that the true mean function is $\varphi(h)$ with $h \in \mathcal{W}_{\alpha,K_0}^\infty$, a Sobolev class with α derivatives, as opposed to $\alpha = 2$.

By deriving a convergence rate such as (16) in this framework, we are able to gain insight on the two important problems in the area of ME: (i) What number of experts m should be chosen, given the size n of the training data? (ii) Given the total number of parameters, is it better to use a few very complex experts, or is it better to combine many simple experts?

For question (i), the results in Theorem 3.3 and Corollary 3.1 suggest that good results can be obtained by choosing the number of experts m to grow as n^r with some

power $r \in (0, 1)$, which may depend on the dimension of the input space and the underlying smoothness of the target function. Smoother target functions and lower dimensions generally encourage us to use less experts.

Question (ii) requires a more detailed study. The complexity of the experts (sub-models) are related to k , the order of the polynomials. We see that increasing k does improve the approximation rate, however this improvement is bounded by the number of derivatives α of the function h . Moreover, this approximation rate is known to be sharp for piecewise polynomials [Windlund, 1977]. The price to pay for this increase in the approximation rate is a larger number of parameters in the model, i.e. a worse estimation error. We will provide below a theoretical result on the optimal choice of k , as well as some numerical evidence.

First of all, an easier expression of the upper bound of the KL divergence in (16) can be derived as $KL \leq O_p(U)$ where $U \equiv (m^{-2(\xi \wedge \alpha)/s} + (m\xi^s)/n)$, where $\xi = k+1$. [This assumes that $v(m) = O(m)$ and uses the fact that the number of parameters needed in s -dimensional polynomials of order k is bounded by $J_k \leq (k+1)^s$.]

We now study the upper bound U , fixing the product $m\xi^s = C$, where C may depend on n and is a bound for the rough order of the total number of parameters.

Proposition 4.1. *Let $\xi = k + 1$ and $U \equiv (m^{-2(\xi \wedge \alpha)/s} + (m\xi^s)/n)$ (\dagger) (which is an upper bound for the KL convergence rate derived in Theorem 3.3). Then the following statements are true:*

I. *Fixing the product $m\xi^s = C$, U is minimized at $\xi = \alpha \wedge (C^{1/s}/e) \equiv \xi_o$. The corresponding optimal $m = \max(e^s, C/\alpha^s)$.*

II. *If α is finite, then U achieves the optimal rate $n^{-2\alpha/(s+2\alpha)}$ under the following choices: ξ is any constant that is at least α and does not vary with n , and $m \approx c_1 n^{s/(s+2\alpha)}$ for any constant $c_1 > 0$.*

III. *If $\alpha = \infty$, the following choices will make U to have a near-parametric rate $U = O((\ln n)^s/n)$: $m \geq 2$ and is constant in n , $\xi \approx c_2 \ln n$ for any constant $c_2 \geq s$.*

Remark 1. *This Proposition suggests that for achieving optimal performance, the ξ (or k , related to the complexity of the experts) and the m (the number of experts) should be compromised. Fixing an upperbound C of the total number of parameters, the optimal $\xi = \alpha \wedge (C^{1/s}/e) \equiv \xi_o$. The optimal compromise therefore depends both on α*

(smoothness of the target function) and s (the dimension of the input space). The formula implies that (a) a smoother target function (indexed by a larger α) will favor more complex submodels (with larger ξ or k), (b) for a very smooth target function (with large enough α), a higher dimension s of the input space will favor the use of simpler submodels (with smaller ξ or $k+1$, possibly smaller than α) and the use of more experts (bigger m).

Remark 2. Although Result I shows how to construct an exactly-optimal compromise between ξ and m , Results II and III show that good convergence rates are quite robust against deviations from these optimal solutions. We note that near-optimal convergence rates can always be achieved with ξ not being too large compared to the sample size n . This is summarized in the two situations in Results II and III, where we see that even in the case $\alpha = \infty$, we only need about $\xi \sim \ln n$ for us to achieve a near-parametric convergence rate.

One drawback of the above theoretic analysis is that it has used a rough upper bound (which has a simple expression) for the total number of parameters associated with k th order s dimensional polynomials. Below we conduct some numeric study, where the exact number of parameters are used. When considering the choice of k , a first impulse is to use polynomials of order $a - 1$, but the number of variables in the model increase exponentially with k if $s > 1$. In fact, in many cases it is preferable to use a smaller k and many experts m if one wishes to control the size of the estimation error. This is consistent with the earlier Remark 1 we made for our theoretic analysis.

Table 1 compares the approximation error using distinct values of k and m holding the estimation error fixed. Assume we have $s = 5$ variables and $\alpha = 6$, a modeler builds a model with $m = 5$ experts and, since it is known that $\alpha = 6$, also chooses $k = 5$. If we further assume $v_m = m \times s$, the total number of parameters in the model is $mJ_k + v_m = 1285$. We can see the smallest approximation error is achieved at $k = 3$ and $m = 21$.

Similarly, fixing the approximation error we see that a balance between m and k is necessary. Fix $\alpha = 6$ and $s = 5$ and assume one wants a model with approximation error proportional to 0.01. Table 2 shows that the model with smaller estimation error that achieves this approximation error is the one with $k = 3$ and $m = 18$.

Table 1: This table compares the approximation error of the model holding the estimation error fixed. We assume $\alpha = 6$ and $s = 5$ and allow for distinct specifications of m and k .

k	m	$m^{-2(k+1)/s}$	$mJ_k + v_m$
0	214	0.1169	1,284
1	117	0.0221	1,287
2	49	0.0094	1,274
3	21	0.0077	1,271
4	10	0.0100	1,310
5	5	0.0210	1,285

Table 2: This table compares the number of parameters of the model holding the approximation error fixed. We assume $\alpha = 6$, $s = 5$ and the estimation error to be proportional to 0.01. We allow for distinct specifications of m and k .

k	m	$m^{-2(k+1)/s}$	$mJ_k + v_m$
0	100,000	0.0100	600,000
1	316	0.0100	3,476
2	46	0.0101	1,196
3	18	0.0098	1,098
4	10	0.0100	1,310
5	7	0.0099	1,799

This quick exercise illustrated one of the main conclusions of this paper: it is not true that one should always use few complex models (small m and large k) or always choose for many complex ones (small k and large m); a balance between k and m should be used instead. Moreover, a small increase in k comparing to the linear model ($k = 1$) can have a good improvement on the approximation and estimation errors.

The results in this paper focus only on target density and mixture-of-experts specified in sections 2.1 and 2.2 respectively. However, similar results can be derived for more complex models and target densities.

5 Conclusion

In this paper we study the mixture-of-experts model with experts in an one-exponential family with mean $\varphi(h_k)$, where h_k is a k^{th} order polynomial and $\varphi(\cdot)$ is the inverse link function. We derive sharp approximation rates with respect to the Kullback-Leibler divergence and convergence rate of the maximum likelihood estimator to densities in an one-parameter exponential family with mean $\varphi(h)$ with $h \in \mathcal{W}_{\alpha,K_0}^\infty$, a Sobolev class with α derivatives. We found that the convergence rate of the maximum likelihood estimator to the true density is $O_p(m^{-2[\alpha \wedge (k+1)]/s} + (mJ_k + v_m)n^{-1} \log n)$, where n is the number of observations, s is the number of covariates, J_k is the number of parameters of the polynomial h_k , m the number of experts and v_m is the number of parameters on the weight functions. Further, if the maximum likelihood estimator is uniquely identified we can remove the “ $\log n$ ” term of the convergence rates.

We discuss model specification and the effects on approximation and estimation errors and conclude that the best error bound is achieved using a balance between k and m , and inclusion of polynomial terms might render better error bounds. Also, the results of this paper can be generalized to more complex target densities and models with simple modifications to the proofs.

We generalize [Jiang and Tanner \[1999a\]](#) in several directions: (i) we assume one can include polynomial terms of the variables on the GLM1 experts; (ii) we assume the target density is in a $\mathcal{W}_{\alpha,K_0}^\infty$ class, for $\alpha > 0$, instead of $\mathcal{W}_{\infty,K_0}^2$; (iii) we show consistency of the quasi-maximum likelihood estimator for fixed number of experts; (iv) we calculate non-parametric convergence rates of the maximum likelihood estimator; (v) we show non-parametric consistency when the number of experts and the sample size increase; and finally (vi) that using polynomials in the experts one can get better estimation and error bounds. These developments have shed light on the important questions of how many experts should be chosen and how complex the experts themselves should be.

Acknowledgements

The authors would like to thank Prof. Marcelo Fernandes and Prof. Marcelo Medeiros for insightful discussion about mixture-of-experts and comments on previous versions of this manuscript.

References

A. Barron and C. Sheu. Approximation of density functions by sequences of exponential families. *The Annals of Statistics*, 19(3):1347–1369, 1991.

C. Bates and H. White. A unified theory of consistent estimation for parametric models. *Econometric Theory*, 1(2):151–178, 1985.

A. Carvalho and M. Tanner. Modeling nonlinear time series with local mixtures of generalized linear models. *Canadian Journal of Statistics*, 33(1), 2005a.

A. Carvalho and M. Tanner. Mixtures-of-experts of autoregressive time series: asymptotic normality and model specification. *IEEE Transactions on Neural Networks*, 16(1):39–56, 2005b.

A. Carvalho and M. Tanner. Modeling nonlinearities with mixtures-of-experts of time series models. *International Journal of Mathematics and Mathematical Sciences*, 9:1–22, 2006.

A. Carvalho and M. Tanner. Modelling nonlinear count time series with local mixtures of poisson autoregressions. *Computational Statistics and Data Analysis*, 51(11):5266–5294, 2007.

G. Celeux, M. Hurn, and C. Robert. Computation and inferential difficulties with mixture distributions. *Journal of the American Statistical Association*, 99:957–970, 2000.

X. Chen. Large sample sieve estimation of semi-nonparametric models. *Handbook of Econometrics*, 6, 2006.

A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.

J. Geweke. Interpretation and inference in mixture models: simple mcmc works. *Computational Statistics and Data Analysis*, 51:3529–3550, 2007.

J. Geweke and M. Keane. Smoothly mixing regressions. *Journal of Econometrics*, 138(1):252–290, 2007.

G. Huerta, W. Jiang, and M. Tanner. Time series modeling via hierarchical mixtures. *Statistica Sinica*, 13(4):1097–1118, 2003.

R. Jacobs, M. Jordan, S. Nowlan, and G. Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.

R. Jennrich. Asymptotic properties of non-linear least squares estimators. *The Annals of Mathematical Statistics*, 40(2):633–643, 1969.

W. Jiang and M. Tanner. Hierarchical mixtures-of-experts for exponential family regression models: approximation and maximum likelihood estimation. *The Annals of Statistics*, pages 987–1011, 1999a.

W. Jiang and M. Tanner. On the identifiability of mixtures-of-experts. *Neural Networks*, 12(9):1253–1258, 1999b.

M. Jordan and R. Jacobs. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2):181–214, 1994.

E. Lehmann. *Theory of point estimation*. Thomson Brooks/Cole, 1991.

P. McCullagh and J. Nelder. *Generalized linear models*. Chapman & Hall/CRC, 1989.

E. Mendes, A. Veiga, and M. Medeiros. Estimation and asymptotic theory for a new class of mixture models. *Manuscript, Pontifical Catholic University of Rio de Janeiro*, 2006.

A. Norets. Approximation of conditional densities by smooth mixtures of regressions. *The Annals of Statistics*, 38(3):1733–1766, 2010.

F. Peng, R. Jacobs, and M. Tanner. Bayesian inference in mixtures-of-experts and hierarchical mixtures-of-experts models with an application to speech recognition. *Journal of the American Statistical Association*, pages 953–960, 1996.

C. Stone. Optimal rates of convergence for nonparametric estimators. *The Annals of Statistics*, 8(6):1348–1360, 1980.

C. Stone. Additive regression and other nonparametric models. *The Annals of Statistics*, 13(2):689–705, 1985.

S. van der Geer. *Empirical Processes in M-Estimation*. Cambridge Univ Pr, 2000.

A. van der Vaart and J. Wellner. *Weak Convergence and Empirical Processes*. Springer-Verlag New York, Inc., 1996.

M. Villani, R. Kohn, and P. Giordani. Regression density estimation using smooth adaptive gaussian mixtures. *Journal of Econometrics*, 153(2):155–173, 2009.

H. White. *Estimation, inference and specification analysis*. Cambridge Univ Pr, 1996.

O. Windlund. On best error bounds for approximation by piecewise polynomial functions. *Numerische Mathematik*, 27:327–338, 1977.

S. Wood, W. Jiang, and M. Tanner. Bayesian mixture of splines for spatially adaptive nonparametric regression. *Biometrika*, 89(3):513–528, 2002.

S. Wood, R. Kohn, R. Cottet, W. Jiang, and M. Tanner. Locally adaptive nonparametric binary regression. *Journal of Computational and Graphical Statistics*, 17(2):352–372, 2008.

S. Wood, O. Rosen, and R. Kohn. Bayesian mixtures of autoregressive models. *Journal of Computational and Graphical Statistics*, 20(1):174–195, 2011.

L. Xu and M. Jordan. On convergence properties of the em algorithm for gaussian mixtures. *Neural computation*, 8(1):129–151, 1996.

Y. Yang and A. Barron. An asymptotic property of model selection criteria. *IEEE Transactions on Information Theory*, 44(1):95–116, 2002. ISSN 0018-9448.

Y. Yang and J. Ma. Asymptotic convergence properties of the em algorithm for mixture of experts. *Neural Computation*, pages 1–29, 2011.

D. Young and D. Hunter. Mixtures of regressions with predictor-dependent mixing proportions. *Computational Statistics & Data Analysis*, 54(10):2253–2266, 2010.

A. Zeevi, R. Meir, and V. Maiorov. Error bounds for functional approximation and estimation using mixtures of experts. *IEEE Transactions on Information Theory*, 44(3):1010–1025, 1998.

A Showing the convergence rate

In this appendix we explain and justify the main steps in proving the convergence rate.

One of the drawbacks of working with the Kullback-Leibler divergence is that it is not bounded. An alternative is to use the Hellinger distance.

Definition A.1 (Hellinger Distance). *Let P and Q denote two probability measures absolute continuous with respect to some measure λ . The Hellinger distance between P and Q is given by*

$$d_h(P, Q) = \left\{ \frac{1}{2} \int \left(\sqrt{\frac{dP}{d\lambda}} - \sqrt{\frac{dQ}{d\lambda}} \right)^2 d\lambda \right\}^{1/2}. \quad (18)$$

Alternatively, the Hellinger distance between two densities p and q with respect to λ is given by

$$d_h(p, q) = \left\{ \frac{1}{2} \int (\sqrt{p} - \sqrt{q})^2 d\lambda \right\}^{1/2}. \quad (19)$$

One can show that if the likelihood ratio is bounded, the KL divergence is bounded by a constant times the square of the Hellinger distance. We use the following result due to [Yang and Barron \[2002\]](#), which is presented together with a basic inequality relating the Hellinger distance and Kullback-Leibler divergence.

Lemma A.1 ([Yang and Barron \[2002\]](#)). *Let $p_{xy} = dP$ and $\|p_{xy}/f\|_{\infty, \Omega \times A} < c_s^2$, for $f \in \mathcal{F}_{m,k}$. Then*

$$d_h^2(p_{xy}, f) \leq KL(p_{xy}, f) \leq 2(1 + \log c_s)d_h^2(p_{xy}, f),$$

where $d_h(p, f)$ stands for the Hellinger distance between the densities p and f with respect to λ .

This Lemma implies that the Kullback-Leibler divergence is bounded by the square of the Hellinger distance, and therefore the convergence rate in the square of the Hellinger distance is the same as the convergence rate in the Kullback-Leibler divergence. The only problem is that, in general, the boundedness condition does not hold on the whole set A (the support of Y). One could overcome this complication by finding the convergence rate inside some subset of A where the KL divergence is bounded and control the tail probability outside this subset.

Let $S(Y, X)$ denote a scalar function of $(Y_1, X'_1), \dots, (Y_n, X_n)'$ and $B(\beta) = \{y \in A : |y| \leq \beta\}$. For every $K \in \mathbb{R}$,

$$P(S(Y, Y) > K) \leq P(\{S(Y, X) > K\} \cap B(\beta)) + P(|Y| > \beta). \quad (20)$$

If A is bounded, we can choose $\beta = \text{ess sup } |A|$, and the second term on the right hand side will be zero. Otherwise, we can take β to be large enough such that $P(|Y| > \beta)$ is small or converges to zero at some rate.

In order to bound the estimation error we shall use results from the theory of empirical processes. The convergence rate theorem presented below is derived for the i.i.d. case, however the same result holds for martingales (see [van der Geer \[2000\]](#)).

To control the estimation rate inside a class of functions we have to measure how big is the class. Let $N_B(\varepsilon, \mathcal{F}, \|\cdot\|)$ denote the number of ε -brackets³ with respect to the distance $\|\cdot\|$, needed to cover the set \mathcal{F} and $H_B(\varepsilon, \mathcal{F}, \|\cdot\|) = \log N_B(\varepsilon, \mathcal{F}, \|\cdot\|)$ the respective *bracketing entropy*. Moreover, let *const.* denote some finite universal constant that may change each time it appears, and write $\mathcal{F}_{m,k}^{1/2} = \{\sqrt{f} : f \in \mathcal{F}_{m,k}\}$. We show that, under some conditions,

$$H_B(\varepsilon, \mathcal{F}_{m,k}^{1/2}, \|\cdot\|_{2,\Omega \times A}) \leq \text{const.} (mJ_k + v_m) \log \frac{C}{\varepsilon},$$

for some finite constant C not depending on ε .

³For a formal definition of *Bracketing Numbers* see [van der Vaart and Wellner \[1996\]](#) chapters 2.1 and 2.7

Hence, our first task is to find the bracketing entropy of $\mathcal{F}_{m,k}^{1/2}$. Assumption 1 implies that

$$\frac{\partial \sqrt{f_{m,k}}}{\partial \zeta'} \leq \frac{\sqrt{f_{m,k}}}{2} [c_g(x)', \delta_1 F(x,y)', \dots, \delta_m F(x,y)'] ,$$

where $0 \leq \delta_i = g_i \pi(h_k(x; \theta_i), y) / f_{m,k} \leq 1$ and $\sum_i \delta_i = 1$.

Hence, for any f_1 and f_2 in some $\mathcal{F}_{m,k}$ indexed respectively by the parameter vectors ζ_1 and ζ_2 in $V_m \times \Theta_{mk}$,

$$|\sqrt{f_1} - \sqrt{f_2}| \leq c(x, y) |\zeta_1 - \zeta_2|_2, \quad (21)$$

with $c(x, y)^2 = (\sqrt{f_1}/2)[|F(x, y)'F(x, y)| + |c_g(x)'c_g(x)|]$. Therefore, the square-root densities in $\mathcal{F}_{m,k}$ are Lipschitz in parameters, with Lipschitz function $c(x, y) \in L^2(\lambda, \Omega \times A)$.

Lemma A.2 (Bracketing Entropy). *Under assumption 1, for any $0 < \delta \leq 1$,*

$$H_B(\varepsilon, \mathcal{F}_{m,k}^{1/2}, \|\cdot\|_2) \leq \text{const.}(mJ_k + v_m) \log \frac{C}{\varepsilon}, \quad (22)$$

where $C = 2\|c(X, Y)\|_{2,\Omega \times A} \text{diam}(V_m \times \Theta_{mk})$; and

$$\int_{0^+}^\delta \log H_B^{1/2}(u, \mathcal{F}_{m,k}^{1/2}, \|\cdot\|_2) du \leq \text{const.}(mJ_k + v_m)^{1/2} \delta \log^{1/2} \frac{C}{\delta} \quad (23)$$

Proof. The first part of the lemma follows from Lemma C.5 and equation (21) together with assumption 1.

The second inequality follows from Lemma C.2. If we take $b = \delta$ and $C > e^\pi$, we have

$$\int_{0^+}^\delta H_B^{1/2}(u, \mathcal{F}_{m,k}^{1/2}, \|\cdot\|_2) du \leq \text{const.}(mJ_k + v_m)^{1/2} \delta \log^{1/2} \frac{C}{\delta}.$$

Proving the lemma. \square

Lemma A.1 requires the likelihood ratio $|p_{xy}/f_{m,k}^*|$ to be bounded. Next lemma shows the rate of decay for the tail probability $P(|Y| > \beta)$, as a function of m and J_k .

Lemma A.3. *Let $p \in \Pi(\mathcal{W}_{\alpha, K_0}^\infty)$ and consider densities on $\mathcal{F}_{m,k}$. Then, under assumption 5, in a set with probability not smaller than $1 - \eta$:*

$$\left\| \sup_p \inf_{f \in \mathcal{F}_{m,k}} \log \frac{p}{f} \right\|_{\infty, \Omega} \leq \text{const.}(mJ_k)^{-\alpha/s} (c\dot{a}_\infty + \dot{b}_\infty) \quad (24)$$

where $\eta = \|\text{Var}(Y|X=x)/c^2\|_{\infty,\Omega}$, c is some large constant, possibly, depending on m and k , and the constants \dot{a}_∞ and \dot{b}_∞ are defined as

$$\begin{aligned}\dot{a}_\infty &= \text{ess sup}_{\Omega \times \Theta} |\partial a(h_k(x, \theta))/\partial \theta|, \quad \text{and} \\ \dot{b}_\infty &= \text{ess sup}_{\Omega \times \Theta} |\partial b(h_k(x, \theta))/\partial \theta|.\end{aligned}$$

Proof. Set $\dot{a}_\infty = \text{ess sup}_{\Omega \times \Theta} |\partial a(h_k(x, \theta))/\partial \theta|$ and $\dot{b}_\infty = \text{ess sup}_{\Omega \times \Theta} |\partial b(h_k(x, \theta))/\partial \theta|$. For ease of notation, also set $h_{ki}(\cdot) = h_k(\cdot, \theta_i)$. By the convexity of the logarithm,

$$\begin{aligned}\log \frac{p}{f} &\leq \sum_{i=1}^m g_i \log(p/\pi_i) \\ &= \sum_{i=1}^m g_i (ya(h(x)) - a(h_{ki}(x))) + b(h(x)) - b(h_{ki}(x)) \\ &\leq (y\dot{a}_\infty + \dot{b}_\infty) \left[\sum_{i=1}^m |g_i - I_{\Omega_i}| |h(x) - h_{ki}(x)| + \sum_{i=1}^m I_{\Omega_i} |h(x) - h_{ki}(x)| \right]\end{aligned}$$

Then, by Assumption 5 and proceeding the same way as in the proof of Theorem 3.1, and taking any value c

$$\text{ess sup}_{x \in \Omega, |y| \leq c} \left| \sup_p \inf_{f \in \mathcal{F}_{m,k}} \log \frac{p}{f} \right| \leq \text{const.} (mJ_k)^{-\alpha/s} (c\dot{a}_\infty + \dot{b}_\infty)$$

The result follows by a simple application of the Chebyschev's inequality. \square

The bound on (24) by itself is not enough since we need to relate the function f^∞ satisfying (24) with $f_{m,k}^*$. It follows from Lemma C.3 that for any $0 \leq c_l < 1$

$$\log \frac{p}{f_{m,k}^*} \leq \frac{1}{(1-c_l)} \log \frac{p}{c_l p + (1-c_l)f_{m,k}^*}.$$

If we choose c_l small enough, we can find a c_p satisfying

$$\log \frac{p}{c_l p + (1-c_l)f_{m,k}^*} \leq c_p \log \frac{p}{f^\infty}.$$

Combining this result with the previous lemma we have inside $B(\beta)$

$$\text{ess sup}_{\Omega \times \Theta} |p_{xy}/f_{m,k}^*| \leq c_\infty \exp \left[\text{const.} (mJ_k)^{-\alpha/s} (c\dot{a}_\infty + \dot{b}_\infty) \right],$$

where $c_\infty = e^{c_p/(1-c_l)}$.

Now we can use theorem 10.13 in [van der Geer \[2000\]](#) to show the rate of convergence of the Hellinger distance between maximum likelihood estimator and the true density. For sake of completeness the theorem is shown below

Theorem A.1 (Theorem 10.13 in [van der Geer \[2000\]](#), pg. 190). *Let $\hat{f}_{m,k}$ denote the maximum likelihood estimator of p_{xy} over $\mathcal{F}_{m,k}$. Set*

$$\bar{\mathcal{F}}_{m,k}^{1/2}(\delta) = \left\{ \sqrt{\frac{f + f^*}{2}} : f \in \mathcal{F}_{m,k}, d_h \left(\frac{f + f^*}{2}, f^* \right) \leq \delta \right\},$$

for some fixed $f^* \in \mathcal{F}_{m,k}$ and let $\|p_{xy}/f^*\|_{\infty, \Omega \times A} \leq c_s^2$. Choose

$$\Psi(\delta) \geq \int_{0^+}^{\delta} \log H_B^{1/2}(u, \bar{\mathcal{F}}_{m,k}^{1/2}(\delta), \|\cdot\|_2) du \vee \delta,$$

in such a way that $\Psi(\delta)/\delta^2$ is a non-increasing function of δ . Then, for $\sqrt{n}\delta_n^2 \geq \text{const. } \Psi(\delta_n)$, we have

$$d_h(p_{xy}, \hat{f}_{m,k}) = O_p(\delta_n + d_h(f_{m,k}^*, p_{xy})).$$

B Proof of the main results

Proof of theorem 2.1. The data generating process of (x, y) and the structure of the model $\{\cdot\}_{m,k}$ is enough to satisfy the measurability assumptions, i.e. it is a weighted sum of measurable functions. also, for any fixed (x_i, y_i) , each π_j is a continuous function of ζ P_{xy} -almost surely, and the same holds for the g'_j 's, then, $f_{m,k}(X_i, Y_i; \cdot)$ is a continuous function of the parameters P_{xy} -a.s. The result follows from theorem 2.12 in [White \[1996\]](#).

□

Proof of Theorem 2.2. There are different approaches to show consistency of the estimate, we proceed by verifying the conditions of theorem 3.5 in [White \[1996\]](#).

The first assumptions regarding the existence of the estimate, are already shown to be satisfied in theorem 2.1. Assumption 3.2, regarding identifiability is satisfied by assumption 3 and 4. It remains to satisfy assumption 3.1, regarding boundedness and uniform convergence of the log-likelihood function. We can show continuity of $\mathbb{E} \log f_{m,k}(X, Y; \zeta)$ by noting that we can interchange integration with limits and a first

order Taylor expansion:

$$\begin{aligned}\mathbb{E} \left[\log \frac{f_{m,k}(X, Y; \zeta)}{f_{m,k}(X, Y; \zeta - \varepsilon)} \right] &\leq \sup \mathbb{E} \left[\left| \varepsilon' \frac{\partial}{\partial \zeta} f_{m,k}(X, Y; \zeta) \right| \right] \\ &\leq \sup \mathbb{E} \left[\left| \frac{\partial}{\partial \zeta} f_{m,k}(X, Y; \zeta)' \frac{\partial}{\partial \zeta} f_{m,k}(X, Y; \zeta) \right| \right]^{1/2} (\varepsilon' \varepsilon)^{1/2},\end{aligned}$$

which is bounded by lemma C.1 and by the fact that ε is arbitrary.

To show uniform convergence of the likelihood function, we satisfy the conditions of theorem 2 in [Jennrich \[1969\]](#). By assumption, $V_m \times \Theta_{m,k}$ is a compact subset of $\mathbb{R}^{v_m \times m J_k}$. Measurability and continuity conditions are already satisfied, then it remains to show that $\log f_{m,k}$ is bounded by an integrable function. Note that we can bound the log-likelihood function by:

$$\begin{aligned}\left| \log \frac{f_{m,k}(X, Y; \zeta)}{\varphi(X, Y)} \right| &= \left| \log \sum_{i=1}^m g_i(X; \nu)(\pi(h_k(X; \theta_i), y) - c(y)) \right| \\ &\leq \sum_i g_i(X; \nu) |[a(h_k(X; \theta_i))Y + b(h_k(X; \theta_i))]| \\ &\leq \max_{1 \leq i \leq m} \text{ess sup}_{x \in \Omega} |a(h_k(x; \theta_i))| |Y| + |b(h_k(x; \theta_i))|\end{aligned}$$

Define the bounding function $D(X, Y) = \max_{1 \leq i \leq m} \text{ess sup}_{x \in \Omega} [|a(h_k(x; \theta_i))| \times |Y| + |b(h_k(x; \theta_i))|]$. The function $|h_k(x; \theta)| \leq \sum_{i=1}^{J_k} |\theta_i| < \infty$ because $\max_i x_i = 1$ and $\sum_i |\theta_i| < \infty$, then both $a(h_k)$ and $b(h_k)$ are finite. Thus, it is straightforward to show that $\mathbb{E} D(X, Y) \leq \infty$, given that $\mathbb{E}_{y|x}(Y) \leq \infty$, which is satisfied by assumption about $p(y|x)$.

Then, $\log f_{m,k}(X, Y; \hat{\zeta}) \rightarrow_{a.s.} \log f_{m,k}(X, Y; \zeta^*)$ as $n \rightarrow \infty$. Therefore, by theorem 3.5 in [White \[1996\]](#), $\hat{\zeta}_n \rightarrow \zeta^*$ P_{XY} -a.s. as $n \rightarrow \infty$.

□

Proof of Theorem 3.1. To bound the approximation rate of the Kullback-Leibler divergence it is enough to bound the upper divergence $\mathcal{D}(f_{m,k}, p)$,

$$\mathcal{D}(f_{m,k}, p) = \int_{\Omega} \sum_{j=1}^{r_m} g_j(x; \nu) \{h_k(x, \theta_j) - h(x)\}^2 dP_x \quad (25)$$

Assumption 5 ensure the existence of a ν_{c_1} such that $\max_j \|g_j(\cdot; \nu_{c_1}) - I_{Q_j^m}(\cdot)\|_{d, P_x} \leq c_1/r_m \|dP_x/d\lambda\|_{\infty, \Omega}$, where $\|dP_x/d\lambda\|_{\infty, \Omega}$ is finite because P_x has continuous density

function with respect to the finite measure λ on Ω . Consider

$$\begin{aligned} \mathcal{D}(f_{m,k}, p) &\leq \underbrace{\left\| \sum_{j=1}^{r_m} \{g_j(\cdot; v_\varepsilon) - I_{Q_j^m}(\cdot)\} \{h_k(\cdot; \theta_j) - h(\cdot)\}^2 \right\|_{1,P_x}}_{(A_1)} \\ &\quad + \underbrace{\left\| \sum_{j=1}^{r_m} I_{Q_j^m}(\cdot) \{h_k(\cdot; \theta_j) - h(\cdot)\}^2 \right\|_{1,P_x}}_{(A_2)}. \end{aligned} \quad (26)$$

Now we just have to find bounds for both terms in the right hand side of (26)(A_1 and A_2). The second term can be written as

$$\begin{aligned} (A_2) &= \int \sum_{j=1}^{r_m} I_{Q_j^m}(\cdot) \{h_k(\cdot; \theta_j) - h(\cdot)\}^2 dP_x \\ &= \int \left\{ \sum_{j=1}^{r_m} I_{Q_j^m}(\cdot) [h_k(\cdot; \theta_j) - h(\cdot)] \right\}^2 dP_x, \end{aligned}$$

where the equality follows from the fact that $I_{Q_j^m} I_{Q_i^m} = I_{Q_j^m} I_{i=j}$, and $\sum_j I_{Q_j^m}(\cdot) = 1$.

If $k < \alpha$, one can choose θ_j such that $\sup_{x \in Q_j^m} |h_k(x; \theta_j) - h(x)| \leq [K_0/(\mathbf{k} + 1)!] \text{diam}(Q_j^m)^{k+1}$ where $\mathbf{k} = (k_1, \dots, k_s)$ is an integer vector satisfying $|\mathbf{k}| = k$. This claim follows from a Taylor expansion of $h(x)$ around fixed points $x_j \in Q_j^m$ and the fact that $h \in \mathcal{W}_{\alpha, K_0}^\infty$. Similarly, if $k \geq \alpha$ we can only use the expansion up to α terms. By assumption 5, $\sup_j \text{diam}(Q_j^m) \leq 1/r_m^{1/s}$. Then

$$\sup_j \sup_{x \in Q_j^m} |h_k(x; \theta_j) - h(x)| \leq \frac{c_0}{r_m^{[\alpha \wedge (k+1)]/s}}, \quad (27)$$

where c_0 depends only on K_0 and $\min(k+1, \alpha)$.

Therefore, $(A_2) \leq c_0^2 / r_m^{2[\alpha \wedge (k+1)]/s}$. Note that

$$\begin{aligned} (A_1) &\leq \sum_{j=1}^{r_m} \|\{g_j(\cdot; v_\varepsilon) - I_{Q_j^m}(\cdot)\} \cdot \{h_k(\cdot; \theta_j) - h(\cdot)\}^2\|_{1,P_x} \\ &\leq \sup_j \sup_{x \in Q_j^m} |h_k(x; \theta_j) - h(x)|^2 \sum_{j=1}^{r_m} \|g_j(\cdot; v_\varepsilon) - I_{Q_j^m}(\cdot)\|_{1,P_x} \\ &\leq \frac{c_0^2 c_1}{r_m^{2[\alpha \wedge (k+1)]/s}}, \end{aligned}$$

where the last inequality is due to equation (27) and assumption 5.

Combining the results for (A_1) and (A_2) ,

$$\mathcal{D}(p, f_{m,k}) \leq \left(\frac{c_0}{r_m^{[\alpha \wedge (k+1)]/s}} \right)^2 (c_1 + 1). \quad (28)$$

It follows from lemma 3.1 that

$$KL(p, f_{m,k}) \leq \left(\frac{c_0}{r_m^{[\alpha \wedge (k+1)]/s}} \right)^2 M_\infty(c_1 + 1). \quad (29)$$

By assumption, $\{r_m\}$ is sub-geometric, then there exists r_m such that $r_m \leq m < r_{m+1}$, and $1/r_m^{2\alpha/s} \geq 1/m^{2\alpha/s} > 1/r_{m+1}^{2\alpha/s}$. Then, $1/(r_m J_k)^{2\alpha/s} \geq 1/(m J_k)^{2\alpha/s} > 1/(r_{m+1} J_k)^{2\alpha/s}$. By definition of $\mathcal{F}_{r_m,k}^* \subseteq \mathcal{F}_{m,k}^* \subset \mathcal{F}_{r_{m+1},k}^*$. Hence,

$$\begin{aligned} \inf_{f_{m,k} \in \mathcal{F}_{m,k}^*} KL(p, f_{m,k}) &\leq \inf_{f_{m,k} \in \mathcal{F}_{r_m,k}^*} KL(p, f_{m,k}) \\ &\leq \frac{M_2 c_2}{r_m^{2[\alpha \wedge (k+1)]/s}} \\ &\leq \frac{c_3}{r_{m+1}^{2[\alpha \wedge (k+1)]/s}} \\ &\leq \frac{c_3}{m^{2[\alpha \wedge (k+1)]/s}}, \end{aligned}$$

where $c_2 = M_\infty^2 c_0^2 (c_1 + 1)$ and $c_3 = M_2 c_2$ does not depend on f . Therefore,

$$\sup_{p \in \Pi(\mathcal{W}_{\alpha,K_0}^\infty)} \inf_{f_{m,k} \in \mathcal{F}_{m,k}^*} KL(p, f_{m,k}) \leq \frac{c_3}{m^{2[\alpha \wedge (k+1)]/s}}$$

□

Proof of Theorem 3.2. The first step is to use equation (20) to bound the convergence rate inside $B(\beta)$ and bound the tail probability. Choose $\beta = c = (m J_k)^{\alpha/\tau}$, and take $c_s^2 = c_\infty e^{const.[\dot{a}_\infty + \dot{b}_\infty]}$. It follows from Lemma A.3 and the discussion afterwards that inside $B(\beta)$

$$\begin{aligned} \|p_{xy}/f_{m,k}^*\|_{\infty,\Omega} &\leq c_\infty \exp\{const.[\dot{a}_\infty + (m J_k)^{-\tau/s} \dot{b}_\infty]\} \\ &\leq c_\infty \exp\{const[\dot{a}_\infty + \dot{b}_\infty]\} \\ &= c_s^2. \end{aligned}$$

This choice of β gives us

$$\eta = P(|Y| > \beta) = \text{ess sup}_x \text{Var}(Y|X=x)(m J_k)^{-2\tau/s}, \quad (30)$$

by lemma A.3, and

$$\text{ess sup}_x \text{Var}(Y|X=x) = \text{ess sup}_x \frac{\ddot{a}(h(x))\dot{b}(h(x)) - \ddot{b}(h(x))\dot{a}(h(x))}{(\dot{a}(h(x)))^3}$$

which is bounded by definition.

Now, inside $B(\beta)$, we can apply Theorem A.1 in the appendix setting $f^* = f_{m,k}^*$. We use Corollary C.1 to bound the bracketing number of $\bar{\mathcal{F}}_{m,k}^{1/2}(\delta)$. By Lemma A.2

$$\int_{0^+}^\delta H_B^{1/2}(u, \bar{\mathcal{F}}_{m,k}^{1/2}(\delta), \|\cdot\|_2) du \leq \text{const.} (mJ_k + v_m)^{1/2} \delta \log^{1/2} \frac{C}{\delta}. \quad (31)$$

Since $C \propto \sqrt{mJ_k + v_m}$, we can choose $\Psi(\delta) \propto (mJ_k)^{1/2} \delta \log^{1/2} \left(\frac{(mJ_k + v_m)^{1/2}}{\delta} \right)$. This choice of function satisfies $\Psi(\delta)/\delta^2$ is non-increasing, and we can take $\delta_n = (mJ_k + v_m)^{1/2} (\sqrt{\log n/n})$. In fact, this choice of δ_n gives us

$$\begin{aligned} \sqrt{n}\delta_n &\geq \text{const.} \Psi(\delta_n) \\ \delta_n &\geq \text{const.} \sqrt{\frac{mJ_k + v_m}{n}} \log^{1/2} \frac{(mJ_k + v_m)^{1/2}}{\delta_n} \\ &= \text{const.} \sqrt{\frac{mJ_k + v_m}{n}} \left(\frac{1}{\sqrt{2}} \log^{1/2} n - \frac{1}{2} \log \frac{\log n}{\sqrt{2}} \right) \\ &= \frac{\text{const.}}{\sqrt{2}} \delta_n - \frac{\text{const.}}{2} \sqrt{\frac{mJ_k + v_m}{n}} \log \frac{\log n}{\sqrt{2}}. \end{aligned}$$

Hence the convergence rate in Hellinger distance is given by:

$$d_h(\hat{f}_{m,k}, p_{xy}) = O_p \left(d_h(f_{m,k}^*, p_{xy}) + (mJ_k + v_m)^{1/2} \sqrt{\frac{\log n}{n}} \right).$$

Our choice of β allows to apply Lemma A.1 to obtain

$$KL(\hat{f}_{m,k}, p_{x,y}) = O_p \left(KL(f_{m,k}^*, p_{x,y}) + (mJ_k + v_m) \frac{\log n}{n} \right).$$

We use Theorem 3.1 to conclude that, inside $B(\beta)$,

$$KL(p_{x,y}, \hat{f}_{m,k}) = O_p \left(\frac{1}{(m)^{2\tau/s}} + (mJ_k + v_m) \frac{\log n}{n} \right). \quad (32)$$

Combining this rate inside $B(\beta)$ with (30), we arrive in our result (14).

We achieve the best rate (15) by taking $m \propto (\log n/n)^{-s/2\tau+s}$ and substituting this rate in (14).

□

Proof of Theorem 3.3. The proof is parallel to 3.2, with just some small changes to lemma A.2. More precisely, since we have a unique $f_{m,k}^*$ for each (m, k) , we can find the bracketing number inside $\mathcal{F}_{m,k}^{1/2}(\delta)$. The argument is the same with the only

difference that inside $\mathcal{F}_{m,k}^{1/2}(\delta)$, $\text{diam}(V_m \times \Theta_{mk}) = \text{const. } \delta$, removing the log-term on the right hand side of (31).

This change allows to choose $\delta_n = \sqrt{(mJ_k + v_m)/n}$, removing the $\log n$ term of the rate.

□

Proof of Proposition 4.1. I. Under the constraint $m\xi^s = C$, $U = C/n + (C\xi^{-s})^{-2(\xi \wedge \alpha)/s}$. Consider two cases: When $\xi > \alpha$, U obviously increases with ξ . So the optimal $\xi_o \leq \alpha$. When $\xi \leq \alpha$, computing the derivative and we know that the function U is minimized at $\xi = (C^{1/s}/e)$. When this point is to the right of α , the function U decreases for all $\xi \leq \alpha$. So we obtain $\xi_o = \alpha$. When $(C^{1/s}/e)$ is to the left of α , the minimum is achieved at $(C^{1/s}/e)$ and $\xi_o = (C^{1/s}/e)$. Combining these we obtain the minimizer $\xi_o = \alpha \wedge (C^{1/s}/e)$.

II. and III. They are straightforward from (†).

□

C Auxiliary Results

In the next lemma, we use the notation $\partial_\theta = \partial/\partial\theta$, $\partial_{\theta\theta'} = \partial^2/\partial\theta\partial\theta'$, $a_j = a(h_k(x; \theta_j))$, $\dot{a}_j = \partial_\theta a_j$, $\ddot{a}_j = \partial_{\theta\theta'} a_j$ and so on.

Lemma C.1. *Let $f \in \mathcal{F}_{m,k}$. Under assumption 1*

- $\mathbb{E}|\log f| \leq \infty$
- $\mathbb{E}|\nabla \log f| \leq \infty$
- $\mathbb{E}|\nabla \log f|_2^2 \leq \infty$
- if we further assume 3 and 4, then $\mathbb{E}|\nabla^2 \log f| \leq \infty$ and is nonsingular at ζ^* .

Proof. This theorem is proved by calculate the derivatives and bounding it.

First note that a_j and b_j are continuous differentiable functions of $h_k(x; \theta_j)$. Since $|h_k(x; \theta_j)| \leq |\theta_j| < \sqrt{J_k}|\theta_j|_2 < \infty$ for any fixed k , then both a_j and b_j are also bounded. The same reasoning can be applied to \dot{a}_j , \dot{b}_j , \ddot{a}_j and \ddot{b}_j . Also, by definition, $\mathbb{E}|y|^p < \infty$ for any $p \geq 0$. Then

$$\mathbb{E} \log f \leq \mathbb{E}[\log \sum_{j=1}^m g_j \pi_j] \leq \mathbb{E}|\max_j [ya_j + b_j + c(y)]| < \infty.$$

Let $\delta_j = g_j e^{ya_j + b_j} p_x / f \leq 1$ and $c^* = \max_j \|\partial_\nu \log g_j\|_{\infty, \Omega}$, then

$$\begin{aligned}\mathbb{E}|\partial_{\theta_j} \log f| &= \mathbb{E}|\delta_j(y\dot{a}_j + \dot{b}_j)x| < \infty, \\ \mathbb{E}|\partial_\nu \log f| &= \mathbb{E}\left|\frac{\sum_j \dot{g}_j e^{ya_j + b_j}}{f}\right| \leq mc^* < \infty.\end{aligned}$$

The same follows for $\mathbb{E}|\partial_\theta \log f|_2^2$, $\mathbb{E}|\partial_\nu \log f|_2^2$ and $\mathbb{E}|\partial_\theta \log f \partial_\nu \log f|$. Let $\dot{c}^* = \|\partial_\nu \log \dot{g}_j\|_{\infty, \Omega}$, and choose any vector α with appropriate dimensions satisfying $\alpha' \alpha = 1$ then

$$\begin{aligned}\mathbb{E}\alpha'|\partial_{\theta_j \theta'_j} \log f|\alpha &= \mathbb{E}\alpha|\delta_j(1 - \delta_j)(y\dot{a}_j + \dot{b}_j)^2 xx' + \delta_j(y\ddot{a}_j + \ddot{b}_j)xx'|\alpha \\ &\leq 0.25\mathbb{E}|y\dot{a}_j + \dot{b}_j|^2 + \mathbb{E} \max_j |y\ddot{a}_j + \ddot{b}_j| < \infty, \\ \mathbb{E}\alpha'|\partial_{\theta_j \theta'_k} \log f|\alpha &= \mathbb{E}\alpha' - \delta_j \delta_k (y\dot{a}_k + \dot{b}_k)(y\dot{a}_j + \dot{b}_j)xx'|\alpha \\ &\leq \mathbb{E}|(y\dot{a}_k + \dot{b}_k)(y\dot{a}_j + \dot{b}_j)| < \infty, \\ \mathbb{E}\alpha'|\partial_{\theta_j \nu'} \log f|\alpha &= \mathbb{E}\alpha'\left|\frac{e^{ya_j + b_j}(y\dot{a}_j + \dot{b}_j)x\dot{g}'p_x}{f}\right|\alpha \\ &\leq \mathbb{E}\alpha'|\delta_j(y\dot{a}_j + \dot{b}_j)x\mathbf{1}_{\mathbf{v}_m}'|\alpha c^* \\ &\leq \mathbb{E}|y\dot{a}_j + \dot{b}_j|c^* < \infty, \\ \mathbb{E}\alpha'|\partial_{\nu \nu'} \log f|\alpha &= \mathbb{E}\alpha'\left|\frac{\sum_j \ddot{g}_j e^{ya_j + b_j} p_x}{f} - \frac{\sum_j \dot{g}_j e^{ya_j + b_j} p_x}{f} \frac{\sum_j \dot{g}'_j e^{ya_j + b_j} p_x}{f}\right|\alpha \\ &\leq c^*|\dot{c}^*| + c^{*2} < \infty.\end{aligned}$$

Since ζ^* is a maximizer of $\mathbb{E} \log f$ over $\mathcal{F}_{m,k}$, $\mathbb{E}|\nabla^2 \log f|$ has to be non-negative definite. Assumption 4 tells us it is also invertible, therefore $\mathbb{E}|\nabla^2 \log f|$ is positive definite. \square

Lemma C.2. For any $0 < a < b \leq 1$ and a positive constant C ,

$$\int_a^b \log^{1/2} \frac{u}{C} du \leq b \left(\sqrt{\pi} \log^{1/2} \frac{C}{b} \right). \quad (33)$$

Proof. For any $0 < a < b \leq 1$,

$$\begin{aligned} \int_a^b \log^{1/2} \frac{u}{C} du &= 2C \int_{\log^{1/2}(C/b)}^{\log^{1/2}(C/a)} v^2 e^{-v^2} dv \\ &\leq 2C \int_{\log(C/b)}^{\infty} t^{1-3/2} e^{-t} dt \\ &= C\Gamma(3/2, \log(C/b)) \\ &\leq b \left(\sqrt{\pi} + \log^{1/2}(C/b) \right), \end{aligned}$$

where the last inequality follows from

$$\begin{aligned} \Gamma(3/2, x) &= \frac{1}{2}\Gamma(1/2, x) + x^{1/2}e^{-x} \\ &= \sqrt{\pi}\Phi(-\sqrt{2x}) + x^{1/2}e^{-x} \\ &\leq (\sqrt{\pi} + x^{1/2})e^{-x}. \end{aligned}$$

□

Lemma C.3. *Let p and q denote two positive densities. For any $0 \leq c_l < 1$,*

$$\log \frac{p}{q} \leq \frac{1}{1 - c_l} \log \frac{p}{c_l p + (1 - c_l)q}. \quad (34)$$

Proof. By the convexity of the logarithm we have

$$\begin{aligned} \log \frac{p}{c_l p + (1 - c_l)q} &= -\log(c_l + (1 - c_l)q/p) \\ &\geq c_l(-\log 1) + (1 - c_l)(-\log q/p) \\ &= (1 - c_l) \log \frac{p}{q} \end{aligned}$$

□

Lemma C.4 (Lemma 4.2 in [van der Geer \[2000\]](#)). *We have, for f_1 , f_2 and some f^* , that*

$$\sqrt{2} d_h \left(\frac{f_1 + f^*}{2}, \frac{f_2 + f^*}{2} \right) \leq d_h(f_1, f_2). \quad (35)$$

This lemma is similar to lemma 4.2 in [van der Geer \[2000\]](#), with the only difference being that we consider an arbitrary f^* and [van der Geer \[2000\]](#) considers f^* to be the true density. The proof remains unchanged.

Corollary C.1. Let $\bar{\mathcal{F}}_{m,k}^{1/2}(\delta)$ be as in theorem A.1, and

$$\mathcal{F}_{m,k}^{1/2}(\delta) = \left\{ \sqrt{f} : f \in \mathcal{F}_{m,k}, d_h(f, f^*) \leq \delta \right\}.$$

We have that $N_B(\varepsilon, \bar{\mathcal{F}}_{m,k}^{1/2}(\sqrt{2}\delta), \|\cdot\|_2) \leq N_B(\varepsilon, \mathcal{F}_{m,k}^{1/2}(\delta), \|\cdot\|_2)$.

Proof. The proof follows from lemma C.4, taking $f_1 = f$ and $f_2 = f^* = f^*$. We have that an $(\sqrt{2}\varepsilon)$ -bracket net for $\bar{\mathcal{F}}_{m,k}^{1/2}(\sqrt{2}\delta)$ is also an ε -bracket net for $\mathcal{F}_{m,k}^{1/2}(\delta)$, all with respect to $\|\cdot\|_2$. \square

The next lemma provides a bound on the bracketing number of functional classes that are Lipschitz in a parameter.

Lemma C.5 (Theorem 2.7.11 in van der Vaart and Wellner [1996]). *Let $\mathcal{F} = \{f_t : t \in T\}$ be a class of functions satisfying*

$$|f_s(x) - f_t(x)| \leq d(s, t) F(x),$$

for some metric d on T , function F on the sample space and every x . Then for any norm $\|\cdot\|$,

$$N_B(2\varepsilon\|F\|, \mathcal{F}, \|\cdot\|) \leq N(\varepsilon, T, d), \quad (36)$$

where $N(\varepsilon, T, d)$ is the ε -covering number of T with respect to the metric d .

It is straightforward to see that if we set $\dim(T) = d$, $c_T = \text{diam}(T)$, and $C = 4\|F\|_2$,

$$N_B(\varepsilon, \mathcal{F}, \|\cdot\|_2) \leq \left(\frac{c_T C}{\varepsilon} \right)^d$$